

NUMERICAL SIMULATION OF FLOW PAST A TAPERED CYLINDER

Dennis C. Jespersen and Creon Levit
NASA Ames Research Center

NUMERICAL SIMULATION OF FLOW PAST A TAPERED CYLINDER

DENNIS C. JESPERSEN* AND CREON LEVIT†

NASA Ames Research Center
Moffett Field, California.

Abstract

We have computed the unsteady three-dimensional low Reynolds number flow past a tapered cylinder. The spanwise variation in natural shedding frequency results in interesting three-dimensional flow phenomena. Our computed hot-wire and spectral data are very similar to experimental results. The computation was done on the Connection Machine, a massively parallel computer; we highlight the capabilities of the Connection Machine for computation and visualization of three-dimensional unsteady flow fields.

1 Introduction

We have two main goals in this paper. The first is to showcase the capability of the Connection Machine, a massively parallel computer, for computation and visualization of three-dimensional unsteady flow fields. We have implemented an algorithm for the unsteady three-dimensional compressible Navier-Stokes equations and will show computed results for low Reynolds number flow past a tapered cylinder. We have also used the Connection Machine for the flow visualization.

Our second goal is to study an interesting three-dimensional flow problem. We have found it fascinating and enjoyable to study the flow field and to speculate about the mechanisms which produce the interesting dynamics.

In the next section we quickly review the key features of the Connection Machine and of the numerical algorithm. Then we present the numerical algorithm and the flow problem to be studied. Finally we show our computed results and discuss flow visualization on the Connection Machine.

2 The Connection Machine

In this section we review the key features of the Connection Machine [5]. The Connection Machine is a single-instruction, multiple data (SIMD) parallel computer. There can be up to 65536 (64K) physical processing elements, and each physical processing element can configure itself so that it appears to consist of more than one “virtual processor”. The number of virtual processors per single physical processor is called the “VP ratio”; this number must be a power of two, and it is generally more efficient to have as high a VP ratio as possible. For example, if we associate one virtual processor with each grid point, if our grid is of size $64 \times 64 \times 32$, and if we are using 8192 (8K) physical processors, then our VP ratio is $64 \cdot 64 \cdot 32 / 8192 = 16$. The VP ratio is limited by the number of words of memory of each physical processor; the machine at NASA/Ames Research Center used for the work reported here had 32K physical processors, each with 2048 32-bit words of memory. Thus if a user’s code requires 100 words of storage per grid point (this counts all user-declared variables, temporary variables generated by the compiler, and system space for flags, etc.), the maximum VP ratio would be 16, the largest power of two less than $2048/100$. Numerical computations are done in 32-bit arithmetic.

For our purposes we can think of the processors as arranged in a three-dimensional grid, one (virtual) processor per grid point, with each processor connected to its Cartesian nearest neighbors. Communication of data is expensive, relative to computation, and communication of data from one processor to another which is “far away” is generally to be avoided; the sole exception to this occurs when the two processors are a power of two distance apart, in which case relatively efficient communication is possible.

The user interface to the Connection Machine is via a front-end computer, which can be either a VAX, a Lisp machine, or a SUN workstation. Most of the work reported in this paper was done via a Lisp machine front end. We are currently moving to a SUN-4 front end; this should yield improved speed, since the faster front end computer can send instructions to the Connection

*Research Scientist, CFD Branch

†Research Scientist, NAS Applied Research Branch

⁰Copyright ©1991 by the American Institute of Aeronautics and Astronautics, Inc. No copyright is asserted in the United States under Title 17, U.S. Code. The U. S. Government has a royalty-free license to exercise all rights under the copyright claimed herein for Governmental purposes. All other rights are reserved by the copyright owner.

Machine more rapidly and can more rapidly perform any calculations that are done in the front end.

The code used for the calculations reported in this paper was developed on a VAX and Lisp machine. We coded in the *Lisp language, an extension of Lisp for parallel computing [11]. We found the Lisp programming environment very useful and helpful in the code development stage.

The code used for the flow visualization was also written in *Lisp, and used the Connection Machine frame buffer to display results. The flow variables were periodically saved as the Navier-Stokes calculation proceeded. Later these saved solutions were used to define velocity vector fields for unsteady particle trace animations. Three-dimensional unsteady flow calculations give rise to very large amounts of data; every saved solution field for a $64 \times 64 \times 32$ calculation took up 2.6 Mbytes of disk space. We used the Connection Machine Datavault, a mass storage device, and saved about 1500 solution fields.

3 Navier-Stokes Algorithm

We implemented an implicit, approximate-factorization central-difference code for the full Navier-Stokes equations. We are solving the laminar flow equations with no thin-layer assumption. The three-dimensional Navier-Stokes equations in generalized curvilinear coordinates can be found in [9]. We can write the equations as

$$Q_t + E_\xi + F_\eta + G_\zeta = \text{viscous terms},$$

where Q is the vector of conserved variables, E , F , G are the inviscid flux terms, and all the viscous terms are on the right-hand side. We discretize the spatial derivatives with central differences and use a general three-point implicit time-stepping method [13]. Our numerical method then looks like

$$\begin{aligned} (I + \frac{\theta \Delta t}{1 + \phi} [\delta_\xi A + \delta_\eta B + \delta_\zeta C]) \Delta Q^n = \\ - \frac{\Delta t}{1 + \phi} (\text{viscous and inviscid flux terms}) \\ + \frac{\phi}{1 + \phi} \Delta Q^{n-1}. \end{aligned}$$

For $\theta = 1$, $\phi = 0$ we have implicit Euler time differencing, for $\theta = 1$, $\phi = 1/2$ we have a three-point second-order time differencing formula. The δ_ξ , δ_η , δ_ζ are central differencing operators. Here $\Delta Q^n = Q^{n+1} - Q^n$; the matrices A , B , C on the left-hand side are evaluated at Q^n . We allow the addition of second and fourth-order artificial dissipation terms to the right-hand side of the equations; for the calculations reported here we used no second-order dissipation, and only a

small amount of fourth-order dissipation. The matrices A , B , and C are the Jacobian matrices of the inviscid flux functions plus additional terms intended to model the viscous fluxes and artificial viscosity terms.

The implicit operator is approximately factored as [13]

$$(I + \frac{\theta \Delta t}{1 + \phi} \delta_\eta A) (I + \frac{\theta \Delta t}{1 + \phi} \delta_\xi B) (I + \frac{\theta \Delta t}{1 + \phi} \delta_\zeta C).$$

In this form the numerical algorithm would require the solution of block banded systems of equations at each time step. For a further reduction in the complexity of the algorithm, we use a diagonal approximation [8], which reduces the block solves to scalar solves. The implicit operator becomes

$$\begin{aligned} T_\eta (I + \frac{\theta \Delta t}{1 + \phi} \delta_\eta \Lambda_\eta) \cdot N \cdot (I + \frac{\theta \Delta t}{1 + \phi} \delta_\xi \Lambda_\xi) P \\ \cdot (I + \frac{\theta \Delta t}{1 + \phi} \delta_\zeta \Lambda_\zeta) T_\zeta^{-1}, \end{aligned}$$

where Λ_η , Λ_ξ , Λ_ζ are diagonal matrices, and T_η , $N = T_\eta^{-1} T_\xi$, $P = T_\xi^{-1} T_\zeta$, and T_ζ are block diagonal matrices with 5×5 blocks. The matrices T_η , T_ξ , T_ζ are the eigenvector matrices of the inviscid flux Jacobians; they are known analytically. Thus the solution process requires the inversion of scalar tridiagonal matrices in the three coordinate directions η , ξ , ζ , and four 5×5 matrix inversions (actually matrix-vector multiplies) at each grid point.

The tridiagonal systems are solved by parallel cyclic reduction [4, 5]. In this algorithm, we have each row of the matrix associated with one processor. In parallel, each processor forms a linear combination of the row above it, its own row, and the row below it, with the coefficients of the combination chosen so that zeros appear in the entries of the linear combination closest to the main diagonal. Then the process repeats, with the processors getting information at distance 2^{n-1} rows away at the n^{th} stage. Just the first step of this algorithm does more arithmetic than Gaussian elimination (Thomas algorithm) does to solve the linear system completely, but the parallel algorithm is ideally suited for the Connection Machine as it finishes in only $1 + \log_2 N$ stages for an $N \times N$ matrix. The last few stages of the algorithm can involve long-distance communication, and this can be a bottleneck. The amount of long-distance communication can be decreased by an "early cutoff" variant of the algorithm [4, 5].

We did some testing of the early cutoff algorithm for our tapered cylinder problem and found savings in wall clock time ranged from 10–20%; we did not use the early cutoff variant in any of the calculations to be reported here. Wall clock times for one step of the code on a $64 \times 64 \times 32$ grid are on the order of 4.6

to 6.5 seconds for a VAX front end, 2.8 to 5.1 seconds for a SUN front end, and 2.7 to 5.2 seconds for a Lisp machine front end; the lower numbers are for a VP ratio of 8 (thus 16K processors), the higher numbers are for a VP ratio of 16 (thus 8K processors). We found that the tridiagonal solver took about 40% of the machine time for the $64 \times 64 \times 32$ problem on the VAX front end and about 50% of the machine time on the SUN and Lisp Machine front ends.

4 Flow Problem

We have numerically computed three-dimensional laminar unsteady low Reynolds number flow past a tapered cylinder. This flow illustrates a case of flow past a body where the natural vortex shedding frequency varies with spanwise location. This variation causes interesting oscillations and “beating” in the flow downstream of the body. The pioneering work of Gaster was reported in [3]. Recent experimental investigations of this flow have been carried out by van Atta and Piccirillo [12] and by Piccirillo [6, 7]. Gaster investigated the flow about a tapered needle, while Piccirillo studied the flow about a tapered cylinder. The parameters in this problem are the Reynolds number and the taper ratio. We normalize the Reynolds number by diameter at mid-span of the cylinder. We define the taper ratio as $(d_0 - d_1)/H$, where d_0 and d_1 are the diameters at the base and at the top of the cylinder, respectively, and H is the height of the cylinder. See Figure 1 for a sketch of the geometry.




Figure 1: Sketch of geometry (taper greatly exaggerated)

We generated grids essentially consisting of “O”-grids about a two-dimensional cylinder stacked in the spanwise direction. We clustered grid points in the

wake region and in the vicinity of the body; typical grid spacing normal to the body was 0.005 (based on a midspan diameter of 1.0). Figure 2 shows a slice of a grid normal to the spanwise direction. The outer boundary is typically 24 mid-span diameters away from the body. Our inflow Mach number was usually 0.2.




Figure 2: Section of grid

The numerical boundary conditions were as follows. At the body surface we imposed the no-slip condition on the velocity components, used linear extrapolation for the pressure, and imposed a zero normal density gradient. At the far field in the streamwise direction we fixed all variables at inflow values at inflow points, and at outflow points we extrapolated all variables from the neighboring interior point. At the far field in the spanwise direction we set the spanwise velocity to zero and imposed a zero normal gradient condition on the other primitive variables.

We define CFL number as

$$\text{CFL} = \Delta t \cdot \max_{j,k,l} [J \{ |U| + c|\nabla\xi|_2, |V| + c|\nabla\eta|_2, |W| + c|\nabla\zeta|_2 \}]_{jkl},$$

where J is the metric Jacobian, $U = \xi_x u + \xi_y v + \xi_z w$, $V = \eta_x u + \eta_y v + \eta_z w$, $W = \zeta_x u + \zeta_y v + \zeta_z w$, $|\nabla\xi|_2 = \sqrt{\xi_x^2 + \xi_y^2 + \xi_z^2}$, and similarly for $|\nabla\eta|_2$ and $|\nabla\zeta|_2$. We usually used a time step of 0.1, which gave a CFL number of approximately 20. We would typically run 50 steps of the flow solver (the solid wall boundary conditions are gradually turned on over the first 30 steps), then introduce a small (order 10^{-4}) nonsymmetric perturbation into the flow field. After several thousand more steps, unsteady vortex shedding becomes apparent. Without the perturbation, many more steps are required before vortex shedding sets up. See Figure 3 for Mach contours at one cross-flow slice at a particular

time. We typically ran 20000 to 50000 or more steps to gather enough data for Fourier analysis.




Figure 3: Typical Mach contours at mid-span

We simulated a hot-wire probe by writing out the crossflow velocity at each time step at selected spanwise stations. The probe was located about 1/4 mid-span diameter aft of the cylinder and about 0.4 diameters to one side of the centerline.

We tested our code by running a nontapered cylinder. For Reynolds number 200 we computed the Strouhal number at several spanwise locations. The Strouhal number is defined as $St = fL/U$ where f is the shedding frequency, L is a length scale, and U is a velocity scale. The shedding frequency was computed by Fourier analysis of the hot-wire signal. At six selected spanwise locations, the Strouhal numbers we computed ranged from 0.204 to 0.205 (when we ran the flow solver using implicit Euler time differencing) and from 0.206 to 0.208 (when we ran the flow solver using a three-point implicit time differencing formula). The Strouhal number for two-dimensional vortex shedding past a circular cylinder at Reynolds number 200 is 0.19, (using the formula from Williamson [14])

$$St = -3.3265/Re + 0.1816 + 1.6 \times 10^{-4} Re. \quad (1)$$

We found a very weak dependence of Strouhal number on spanwise location, which may be attributable to the manner in which the boundary conditions at the top and bottom of the tapered cylinder were implemented. Spectral analysis found no indication of a low-frequency component in the hot-wire signal, in contrast to two-dimensional calculations where careful choice of far-field boundary conditions may be necessary to avoid spurious wave reflections [1, 2].

We also computed the flow past a nontapered cylinder at Reynolds number 122.5. Spectral analysis of our simulated hot-wire data gave Strouhal number 0.185 at

all spanwise locations. The Strouhal number for two-dimensional vortex shedding past a circular cylinder at Reynolds number 122.5 is 0.174 (from equation 1). Again, there was no sign of spurious numerical wave reflections.

We studied the effect of various numerical parameters on the computed flow past a cylinder with taper ratio 0.01 and Reynolds number ranging from 100 to 145 (Reynolds number 122.5 at mid-span). For time differencing we used either implicit Euler or the three-point implicit method with $\theta = 1$, $\phi = 1/2$; we used meshes of size $64 \times 64 \times 32$, $64 \times 64 \times 64$, and $64 \times 64 \times 128$; we used time steps $\Delta t \in \{0.1, 0.2, 0.4\}$; and the outer boundary location was taken at 24 or 48 mid-span diameters. The distance from the body to the first grid line in the normal direction was taken to be 0.005 mid-span diameters. In all these cases the general features of the flow were the same, at least judging by the hot-wire data; there were some quantitative differences in the spectral peaks with $\Delta t = 0.2$ or 0.4 . We conclude that $\Delta t = 0.1$, outer boundary at 24 mid-span diameters, and 32 spanwise points are sufficient to produce accurate signals for spectral analysis; we also used this case for flow visualization.

We will show some computed results for a cylinder with taper ratio 0.01 and Reynolds number range of 90 to 145, the same range as in the experiment of Piccirillo [6, 7]. In this particular case we used a $64 \times 64 \times 64$ mesh and three-point implicit time differencing; the Reynolds number at mid-span is 117.5. A portion of the simulated hot-wire data is shown in Figure 4 (at the end of the paper). This output is very similar to experimental hot-wire signals in Gaster [3], van Atta and Piccirillo [12] and Piccirillo [6, 7]. Spectral analysis of the hot-wire data gives the results in Figure 5. This analysis uses the value of the crossflow velocity component at 70000 time steps at each of 15 selected spanwise stations. One can see the strictly periodic nature of the signal at spanwise station 4. A much lower frequency modulation is just visible at spanwise station 8 and grows in strength as one moves along the span toward the thinner end of the tapered cylinder. When spanwise location 32 is reached, the originally dominant peak at a nondimensional frequency of about 0.16 has lost dominance; a new dominant peak at nondimensional frequency of about 0.175 has replaced it. This new value of dominant frequency is the sum of the originally dominant frequency and the low frequency modulation visible on the curves at spanwise stations 8 through 28. After spanwise station 32, a peak at nondimensional frequency of about 0.19 gains strength and becomes dominant by about $z = 44$ or $z = 48$. At $z = 52$ this peak is still evident but by $z = 56$ the power spectrum is of broad band type.

The general features of the spectral analysis for all our numerical experiments are as follows: single strong peak at frequency ν_0 for hot-wire station near the base of the cylinder (periodic vortex shedding); emerging second peak at frequency $\nu_1 > \nu_0$ and small secondary peak at frequency $\nu_1 - \nu_0$ as we move up along the cylinder; eventually the two frequencies ν_0 and ν_1 have equal energy, then frequency ν_1 becomes dominant; then (depending on spanwise resolution) a third frequency $\nu_2 > \nu_1$ may emerge as the dominant frequency; and finally a regime where the spectral analysis shows a broad peak and it is difficult to single out a dominant frequency.

One of our other numerical experiments was a cylinder with taper ratio 0.01 and a Reynolds number range of 50 to 150. There was no change in the qualitative behavior of the solution. Another case we ran was a cylinder with taper ratio 0.001 and Reynolds number ranging from 105 to 115. In this case we observed one dominant frequency ν_0 across the entire span of the cylinder, with a second frequency ν_1 beginning to emerge toward the thin end of the tapered cylinder.

5 Flow Visualization

The three-dimensional unsteady flow visualization was carried out in three stages. First the full solution at every tenth time step was saved to the attached disk on the Connection Machine. This resulted in about 3.5×10^9 bytes of saved solution data for a typical long run. Next, a program was run that generated streaksurfaces by advecting a large number of “bubbles” (ideal infinitesimal fluid elements) through the changing velocity field. The bubbles were emitted from “bubble wires” whose location was chosen at the beginning of this advection stage.

For most of our visualizations, two bubble wires were used. Each wire emitted 128 new bubbles at every time step. One wire ran along one side of the cylinder body, near the “left” separation line, and another ran along the other side of the cylinder body, near the “right” separation line¹. For studies of the large scale flow structure, the bubble wires ran almost the whole length of the cylinder, whereas for the detailed studies of the vortex “dislocation” phenomena, the bubble wires ran only about one sixth the length of the cylinder.

The Connection Machine implementation of the advection algorithm is interesting. It makes use of two “virtual processor sets” (VP sets). VP set number one

¹Actually, the instantaneous separation “lines” on the tapered cylinder are not lines at all, but gently waving curves that move with time. However, they do not move very far, and the bubble wire positions were fixed just downstream of the separation lines’ extreme location

typically has 128K virtual processors – one VP per grid-point – and holds velocity field data for the two solution time steps involved in one advection step. A second order Runge-Kutta algorithm with a fixed time step is used for advection with all positions and velocities maintained in computational coordinates. VP set number two consists of one VP per bubble. The number of active processors in this VP set grows as new bubbles are emitted, and shrinks when particles travel off the edge of the computational grid. A maximum of 32K bubbles were advected at each time step, with the very oldest being removed from the flow when necessary in order to be recycled.

At every step of the advection, interprocessor communication between the two VP sets is required. Each bubble (in VP set 2) must get the velocity vectors from the eight corners of the grid cell (in VP set 1) that it is currently passing through. Trilinear interpolation is used to define the velocity at each bubble’s position. This happens twice per step, once for each substep of the Runge-Kutta time advance. Once the bubble’s new computational space position is known, it gathers its physical position information from VP set 1 in a similar manner, and uses trilinear interpolation to compute its physical (x, y, z) location. The updated particle positions are visible on the CM frame buffer during this procedure, and the 3D view can be transformed using the mouse in real time. The physical locations of all currently active particles are also appended to a disk file at the end of every advection step. Typically, after about one hour and 1000 advection steps, the file contains about 4×10^8 bytes.

The final step of the visualization process involves reading some large segment of this file (typically several hundred advection steps worth) into the Connection Machine at one time. Then, as time advances, the three dimensional positions of all the bubbles active at any time step are transformed via a typical perspective projection and displayed on the CM frame buffer. Bubbles can be colored according to the wire that emitted them, or by their age. The frame-rate while playing back these three dimensional streaksurface movies is around 20 frames per second. Position, scale, and 3D rotations are controlled by the mouse in real time. The flow movie can be stopped, rotated, and continued on demand.

While the raw floating point performance of the Connection Machine during the advection phase is rather poor, the complete system can produce a richly detailed unsteady three dimensional flow visualization running at 20 frames per second with the three dimensional view under real time control, from several gigabytes of data, in about an hour. See [10] for more information about flow visualization on the Connection Machine.

Flow visualization of the computed solution fields for unsteady flow past a tapered cylinder reveals unsteady vortex shedding with the vortices being shed at an angle (not parallel) to the cylinder. The visualization shows events that are consistent with the “dislocation” events discussed in [14]. In Figure 6 we show a sequence of frames from the flow visualization. The cylinder is on the left, with the thicker end at the bottom, and flow is from left to right. The portion of the cylinder from spanwise station 18 to spanwise station 24 is shown. Particles are being released from only one side of the cylinder. Reading down the figure, first the left half and then the right, we see the oblique shedding and then the dislocation event.

6 Discussion

Let us first make a comment on the fact, observed both experimentally and computationally, that the shedding frequency increases as we move from the thick end of the cylinder toward the thin end, i.e., as the local Reynolds number decreases. This may seem counter-intuitive, since we know that Strouhal number decreases as the Reynolds number decreases. But consider, from the definitions $St = fL/U$ and $Re = UL/\nu$, and from the experimental determination that $St \approx (A - B/Re)$, we can derive

$$f = USt/L = U^2St/(\nu Re) \approx (U^2/\nu Re)(A - B/Re).$$

With the experimentally determined values of $A = 0.215$ and $B = 5.1064$ ([14]), we can easily see that f decreases as Re increases as long as $Re > 2B/A \approx 47.0$. Thus the shedding frequency decreases as the Reynolds number increases, at least for the range of Reynolds numbers we are considering.

The spectral analysis of our computed results, as exemplified in Figure 5, is consistent with the results of [6] and [7]. The flow field near the cylinder may be thought of as being partitioned into spanwise “cells”, where a cell is defined purely in spectral terms as that spanwise region where a certain frequency is dominant. Thus in Figure 5, one cell extends from the bottom of the cylinder up to about spanwise station 32. At spanwise station 32 a different (higher) frequency has become dominant, and this frequency retains its dominance up to about spanwise station 44. At spanwise station 48 a third higher frequency has become dominant; this frequency remains dominant up to about spanwise station 56, after which the spectral peak is quite broad and no single frequency is dominant.

We can attempt to define a Strouhal number for this flow problem using the spectral analysis. Given a spectral analysis $P(f)$ vs. f of power vs. frequency, we can

Generalized Strouhal number			
Span station	Local Re	St*	St (2D)
4	137.2282	0.1856	0.1793
8	134.3586	0.1818	0.1783
12	131.4890	0.1779	0.1773
16	128.6195	0.1740	0.1763
20	125.7499	0.1704	0.1752
24	122.8804	0.1671	0.1741
28	120.0108	0.1656	0.1730
32	117.1413	0.1699	0.1719
36	114.2717	0.1702	0.1707
40	111.4021	0.1671	0.1695
44	108.5326	0.1666	0.1683
48	105.6630	0.1674	0.1670
52	102.7935	0.1660	0.1656
56	99.9239	0.1643	0.1642
60	97.0543	0.1657	0.1628

Table 1: Generalized Strouhal number calculation for 3D tapered cylinder.

define a “generalized Strouhal number” St^* via

$$St^* = (L/U) \int_0^\nu f P(f) df \bigg/ \int_0^\nu P(f) df,$$

where L is a local length scale, and the upper limit of integration is chosen (by hand) so that the region from 0 to ν contains the first major peak in the spectrum. If we carry out this calculation for the data shown in Figure 5 we obtain Table 1 (the frequency data in Figure 5 are already normalized by $1/U$).

In Table 1 the Strouhal number for two-dimensional cylinder shedding is computed from the empirical formula of equation 1. In all cases the difference between the “generalized Strouhal number” St^* and the Strouhal number from two-dimensional experiments is less than a few percent. This data and data from our other computational experiments are consistent with the hypothesis that the flow in the wake of a tapered cylinder close to the body has its shedding frequency determined by the two-dimensional Strouhal number.

References

- [1] S. Abarbanel, W. Don, D. Gottlieb, D. Rudy, J. Townsend, Secondary Frequencies in the Wake of a Circular Cylinder with Vortex Shedding, ICASE Report 90-16, (1990).
- [2] W.-S. Don and D. Gottlieb, Spectral Simulation of Unsteady Compressible Flow Past a Circular Cylinder, ICASE Report No. 90-29 (1990).

- [3] M. Gaster, Vortex Shedding from Circular Cylinders at Low Reynolds Numbers, *J. Fluid Mech.* **38**, pp. 565—576 (1969).
- [4] R. W. Hockney and C. R. Jesshope, *Parallel Computers 2*, Adam Hilger, 1988.
- [5] D. C. Jespersen and C. Levit, A Computational Fluid Dynamics Algorithm on a Massively Parallel Computer, *Int. J. Supercomputer Applications*, vol. 3, no. 4, pp. 9—27 (1989).
- [6] P. S. Piccirillo, An Experimental Study of Flow Behind a Linearly Tapered Cylinder at Low Reynolds Number, Master's Thesis, University of California at San Diego, 1990.
- [7] P. S. Piccirillo and C. W. Van Atta, An Experimental Study of Flow Behind a Uniformly Tapered Cylinder, AIAA Paper 91-0750, AIAA 29th Aerospace Sciences Meeting, Reno, 1991.
- [8] T. H. Pulliam and D. S. Chaussee, A Diagonal Form of an Implicit Approximate-Factorization Algorithm, *J. Comp. Phys.* **39**, pp. 347—363 (1981).
- [9] T. H. Pulliam and J. L. Steger, Implicit Finite-Difference Simulations of Three-Dimensional Compressible Flow, *AIAA J.* **18**, pp. 159—167 (1978).
- [10] J. A. Sethian, J. B. Salem, and A. F. Ghoniem, Interactive Scientific Visualization and Display Techniques, Thinking Machines Corporation Technical Report VZ88-1, Cambridge, Massachusetts (1988).
- [11] **Lisp Reference Manual*, version 5.0, Thinking Machines Corp., Cambridge, Massachusetts (1987).
- [12] C. W. Van Atta and P. Piccirillo, Topological Defects in Vortex Streets Behind Tapered Circular Cylinders at Low Reynolds Numbers, *New Trends in Nonlinear Dynamics and Pattern Forming Phenomena: The Geometry of Nonequilibrium*, ed. P. Huerre and P. Couillet (1989).
- [13] R. F. Warming and R. M. Beam, On the Construction and Application of Implicit Factored Schemes for Conservation Laws, in *Computational Fluid Dynamics, SIAM-AMS Proceedings Vol. 11*, ed. Herbert B. Keller, American Mathematical Society, Providence, Rhode Island, 1978.
- [14] C. H. K. Williamson, Oblique and Parallel Modes of Vortex Shedding in the Wake of a Circular Cylinder at Low Reynolds Numbers, *J. Fluid Mech.* **206**, pp. 579—627 (1989).




Figure 4: Hot-wire signals (actually cross-flow velocity component)




Figure 5: Spectral analysis of simulated hot-wire, from $z = 4$ (thick end of cylinder) to $z = 60$ (thin end of cylinder)




Figure 6: Flow visualization. Particles are released from one side of cylinder. Reading from top to bottom and left to right, the frames are at nondimensional times $t = 0, 50, 100, 150, 175, 185, 195, 215$. During this time eight vortices are shed from the top of the section and seven from the bottom. At the beginning and end of this sequence, vortices are being shed coherently across this section of the cylinder. The intermediate frames illustrate the vortex dislocation process.